IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent Application of	
J.Thomas Ngo et al.	Group Art Unit: 2177
Application No.: 10/059,233	Examiner: KUEN S. LU
Filed: January 31, 2002	Confirmation No.: 6547
For: DATA REPLICATION BASED UPON A NON-DESTRUCTIVE DATA MODEL	

REQUEST FOR RECONSIDERATION

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

In response to the Office Action dated July 14, 2004, Applicants respectfully request reconsideration and withdrawal of the rejections of the claims.

Claims 1, 2, 5-12, 31-36, 38-76, 84-86, 97-110 and 113-123 were rejected under 35 U.S.C. §103, on the grounds that they were considered to be unpatentable over the Shih et al patent (U.S. 6,615,223) in view of the Neufeld et al patent publication (U.S. 2003/0126391). The remaining claims were rejected as being unpatentable over these two references, in view of various tertiary references. For the reasons presented below, Applicants respectfully traverse these rejections.

The present invention is directed to the replication of data in an environment where different users can edit a data object independently of one another, as a result of which multiple conflicting versions of the data object can exist at one time. The data object might be, for example, a word processing document on which the users are collaborating, or a

data base record. As discussed in the background portion of the application, in the past it was conventional to employ techniques for conflict avoidance or conflict resolution in these types of situations. Systems that employ these techniques are based on a model that assumes that a single authoritative version of the data object should exist at any given time.

In contrast to these approaches, the present invention employs a non-destructive data model, that permits different versions of a data object to coexist. In the context of this approach, the history of changes to a data object are captured within an atom graph. For instance, if two users, A and B, are collaborating on a document, each user might independently make changes to the document on his or her computer. In the context of the invention, a first atom graph reflecting the changes made by user A resides on user A's computer, and a second atom graph reflecting the changes made by user B resides on user B's computer. Subsequently, when the users synchronize their versions of the document, each atom graph is updated by transferring to it the atoms from the other graph that are not already present. In one embodiment of the invention, this transfer is accomplished by forming the mathematical union of the atom graphs at the respective users' computers. See, for example, Figure 3 and the description beginning at page 10, line 26. In essence, the synchronization of the two versions of the document is carried out in a non-destructive manner. Each user's modifications are maintained intact, rather than having one set of modifications overwrite the other in accordance with an established rule.

Claim 1 recites a method for replicating data at multiple devices, which includes the step of representing the history of a data object at each of the devices by means of a graph of atoms in a store, where the graph includes at least one parent atom of a first type that contains information pertaining to an operation performed on the data object, and at least

two descendent atoms of this first type. The next step of claim 1 comprises adding an atom of the first type to the atom graph in the store at a given device when an operation is performed on the data object at that device. In connection with this claimed subject matter, the Office Action refers to the Shih patent at Figures 5 and 11, particularly elements 96, 98, 102 and 104.

The Office Action is apparently interpreting the directory tree structure shown in these figures as an atom graph, and the individual elements 96, 98, 102 and 104 as atoms. It is respectfully submitted, however, that the elements of the tree structure shown in Figures 5 and 11 of the Shih patent do not constitute a graph of atoms that represent the history of a data object. In the context of the claimed invention, an atom graph pertains to a given object, e.g., a word processing document or a record in a database. Each atom in the graph pertains to an operation that is performed on *that* object, e.g., modification or deletion of the object. In contrast, the elements depicted in the directory information tree of the Shih patent do not all pertain to a particular data object. Rather, they represent *different* objects or classes of objects. For example, element 102 represents one record in the directory, pertaining to a person named Joe Jones. Element 104 does not pertain to the same record. Rather, it represents a different record in the directory, that is associated with another person named Bob Last.

Accordingly, even if the individual elements 96-104 can be considered to be "atoms," and the directory tree structure is labeled an atom graph, it is respectfully submitted that this atom graph is not the same as that recited in claim 1. The directory tree does not represent the history of a particular data object. Rather, it identifies the relationship of a number of different data objects contained within the directory. There is

no disclosure in the Shih patent suggesting that a given element, e.g., 102 or 104, identifies an *operation* that is performed on a data object. It only represents the data object, per se, i.e. its current state.

Apparently recognizing this distinction, the Office Action refers to the Neufeld publication. This publication is directed to an objective that is entirely distinct from the field of data replication and synchronization. In particular, the Neufeld publication is concerned with information media that has limited write cycle capability. For example, in paragraph 0008, the publication explains that optical media such as CD-RW only has the ability to withstand a limited number of rewrite cycles.

To overcome this constraint, the Neufeld publication discloses a technique in which the location of a file is moved to different storage areas on the medium in concert with the write cycle capabilities of the medium. Referring to Figure 7A, the first version of a file is stored at location 712. The file stored at this location undergoes a number of rewrite operations, until a predetermined rewrite cycle threshold is reached. Thereafter, when the next rewrite is to occur, the file is written to a new location 714. At the same time, a pointer 704 identifies the new location for the file, and the previous version at location 712 is marked with a symbol 702, to identify the fact that it is unstable, and no longer in use. In a similar manner, when the rewrite cycle threshold is reached at location 714, the most recent version of the file is then moved to location 716, as depicted in Figure 7C.

The Office Action characterizes the operation depicted in Figures 7A-7C as a representation of the history of a data object in a graph of atoms. It is respectfully submitted, however, that the various versions stored at the locations 712, 714 and 716 do not constitute atoms that represent the history of operations performed on a data object.

Rather, each version is capable of being overwritten a number of times, up to the rewrite cycle threshold, and represents only a snapshot of the file at the time that the threshold was reached. For instance, if the rewrite threshold is 100 cycles, version 1 stored at location 712 would only indicate the status of the file after the 100th rewrite. Assume that users A and B both contributed to that version. It may be the case that a modification made by user A was subsequently overwritten by user B. The version stored at location 712 would not provide any information regarding the modification made by user A that was subsequently overwritten. In other words, it does not represent the *history* of the data object. Rather, each of the versions at locations 712, 714 and 716 only show the state of the file at the threshold points, e.g. after every 100th rewrite, but no information regarding the operations that took place intermediate those points.

In the non-destructive data model of the present invention, in which the history of a data object is captured in the atom graph and the respective atom graphs are updated with each other's atoms, it is possible to ascertain the contributions made by A at any given time, even if they did not end up in the final version. The approach taken by the Nuefeld patent does not provide the same result. Since a file can be overwritten many times while the active version is stored in a given location, up to the rewrite threshold, it is not possible to determine the operations that were made but did not end up in the final version.

Accordingly, it is respectfully submitted that the Neufeld publication does not disclose a graph of atoms that represents the history of a data object.

Furthermore, it is respectfully submitted that the Shih and Neufeld references are directed to entirely disparate objectives, and there is no apparent reason for a person of ordinary skill in the art to apply the teachings of one to the other. The Shih patent is

directed to data replication, and more particularly to the problems that are encountered with data replication in a heterogeneous environment, in which the various replication sites have different schema organizations for the data. See, for example, column 2, lines 16-36. To this end, the patent discloses a schema and format-independent method for data replication. In contrast, the Neufeld publication is directed to the limited write cycle capabilities of certain storage media. This publication teaches a technique for extending the useful life of such media, by migrating stored data from one portion of the media to another portion thereof. It is respectfully submitted that these two references have nothing to do with one another, and there is no apparent motivation to apply the teachings of one to the disclosure of the other. Any possible reason for doing so can only be based upon hindsight knowledge of the presently claimed invention.

In proffering the combination of the two references, the Office Action refers to the Shih patent's teaching of a garbage collector daemon process, and appears to be suggesting that it would be obvious to eliminate this process in view of the teachings of the Neufeld publication. It is respectfully submitted, however, that the Neufeld publication contains no teachings that pertain to the preservation of content and operation history on data objects, or that would otherwise suggest the elimination of the garbage collector process. Rather, the Neufeld publication is concerned with a *physical* phenomenon, namely the extension of the useful life of storage media having limited write cycle capabilities. It is respectfully submitted that it would not be obvious to eliminate the garbage collector process disclosed in the Shih patent in light of any of the teachings of the Neufeld publication.

For at least the foregoing reasons, therefore, it is respectfully submitted that the subject matter of the pending claims is not suggested by the Shih and Neufeld publications, whether considered individually or in combination.

Claims 2, 110 and 115 recite that the updating of the history of a data object is carried out by the mathematical union of atom graphs in two devices. In connection with this claimed subject matter, the Office Action refers to Figures 5 and 11 of the Shih patent, particularly the addition of directory entry 104 to the directory tree structure. It is respectfully submitted that this disclosure does not constitute a teaching of forming the mathematical union of two atom graphs. The Shih patent only discloses a single "atom graph," namely the directory tree structure of Figure 5. There is no second directory tree with which it is united. Rather, Figure 11 merely discloses the addition of another entry in the *same* directory. The patent does not disclose a graph union operation that is performed on that structure.

For this additional reason, it is respectfully submitted that the subject matter of claims 2, 110 and 115 is not suggested by the references.

Further distinguishing features of the invention are set forth in the other pending claims. However, in light of the fundamental distinctions set forth above, it is believed that a detailed discussion of these other differences is unnecessary at this time. Furthermore, it is respectfully submitted that the tertiary references that are applied in the rejections of some of the dependent claims do not suggest the foregoing distinctions.

Reconsideration and withdrawal of the rejections, and allowance of all pending claims are respectfully requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

Date: November 10, 2004

By: _____
James A. LaBarre
Registration No. 28,632

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620